

REAL TIME CONTROL PROTOCOL SESSION MATCHING

FIELD OF THE INVENTION

5 The present invention relates generally to network media streams and particularly to monitoring network media streams.

BACKGROUND OF THE INVENTION

Sub A'7
10 Real Time Protocol ("RTP") is the standard protocol defining the real-time transmission of media streams (*e.g.*, voice) over data networks, such as in Voice Over IP. A companion protocol to RTP is the Real Time Control Protocol or RTCP. Referring to Fig. 1, media packets transmitted between A 100 and B 104 and vice versa during a session are formatted and transmitted (continuously) over network 108 according to RTP while additional performance information governing the communication link (*e.g.*, key statistics
15 about the media packets being sent and received by each endpoint (A or B) such as jitter, packet loss, round-trip time, *etc.*) are transmitted (discontinuously) over the network 108 according to RTCP. Endpoints A and B are typically computational components but can be or include any other form of audio or video communications interface. RTCP performance information is useful not only for the session participants, A and B, but also for a network
20 monitor 112. Network administrators can use such information not only for network administration but also for network troubleshooting and management.

 RTCP is specifically designed to provide such information to the network monitor 112 via an IP multicast architecture. In IP multicast, a single packet is transmitted to a group of recipients by first being sent to a multicast address and then being distributed by the network

to multiple addresses associated with the multicast address. The group of recipients includes not only the other (receiving) party in the session, namely A or B as appropriate, but also the monitor. When IP multicast is used for RTCP information, it also should be used for RTP information. Timing calculations for RTP packets are made based on the behavior of RTCP packets, and such timing calculations are only an accurate predictor of RTP packet transmission characteristics if both RTP and RTCP packets are transmitted by IP multicast techniques. IP multicast is generally disfavored because multicast is complicated which causes administration complications and difficulties.

A common architecture for transmitting media streams between two session participants is known as unicast. In unicast, the packets are transmitted to only one and not multiple destinations. It is more difficult for the monitor to collect and analyze RTCP packets as the packets are transmitted only to the other session participant and not to the monitor.

To enable the monitor to obtain RTCP packets, a dual unicast architecture has been developed. In dual unicast, one session participant (A) transmits both RTP and RTCP packets to the other session participant (B) and RTCP packets to the monitor. Dual unicast, however, exposes design limitations in the RTCP protocol itself. Although endpoint session ids are unique to a particular (first) session (such as between A and B), an endpoint in a concurrent (second) session (such as between C and D) can have the same session id or synchronization source id ("SSRC") as an endpoint (A or B) in the other (first) session. When duplicate endpoint session ids are concurrently in use, the monitor can have substantial difficulty determining which RTCP packets correspond to which session, potentially causing inaccurate performance analysis.

By way of illustration, in the example above assume that A sends an RTCP packet addressed to B and an RTCP packet addressed to the monitor. The RTCP packet addressed to the monitor includes A's transport address, A's SSRC, and B's SSRC but does not include the transport address of B. Likewise, C sends an RTCP packet addressed to D and an RTCP packet addressed to the monitor. The RTCP packet addressed to the monitor includes C's transport address, C's SSRC, and D's SSRC but does not include the transport address of D. If B and D have the same SSRC, the monitor is unable to definitively determine that a selected RTCP packet sent to either B or D corresponds to the A-B session or the C-D session.

SUMMARY OF THE INVENTION

These and other needs are addressed by the various embodiments and configurations of the present invention. The present invention generally matches or associates session packets communicated in a session between two or more endpoints or participants with the identities of the participants, *e.g.*, session ids (*e.g.*, SSRC) and/or network addresses (*e.g.*, transport addresses), creating new sessions if appropriate. Each session participant is typically identified by network address (*e.g.*, UDP port and Internet Protocol address) and/or session (*e.g.*, SSRC) id.

In one embodiment, a method for identifying a corresponding session for a packet is provided that includes the steps of:

(a) receiving at least a first packet communicated between first and second endpoints (or participants) to a first session, the first packet comprising at least one of a network

address of the first endpoint (e.g., port or UDP), a session id of the first endpoint (e.g. SSRC), and a session id of the second endpoint;

(b) comparing the at least one of a network address of the first endpoint, a session id of the first endpoint, and a session id of the second endpoint in the packet with a listing of at least one of network addresses and session ids contained in previously received packets; and

(c) when the at least one of a network address of the first endpoint, a session id of the first endpoint, and a session id of the second endpoint in the packet matches an entry in the listing, determining a network address of the second endpoint in the first session (which is typically unknown).

The listing can be a single table or a plurality of tables. In one configuration, the listing is in the form of or includes an active session table, which includes the network addresses of all of the participants of a selected session and optionally the session ids of all of the selected session participants. In one configuration, the listing is in the form of or includes an orphan session table, which includes, in each entry, the network address of one (but not the other) session participant and optionally the session ids of one or both participants. The active and orphan session tables are typically disjoint.

In another embodiment, a method for monitoring a multi-party session is provided that includes the steps of:

(a) receiving, at a first endpoint, at least a first packet communicated between the first endpoint and a second endpoint to a first session, the first packet comprising a network address of the first endpoint and a network address of the second endpoint; and

(b) transmitting at least a second packet to a session monitor, the second packet including the respective first and second network addresses of the first and second endpoints. In one configuration, the first packet is effectively retransmitted or reflected by the receiving first endpoint to the session monitor. This is particularly useful to third-party endpoints that do not implement dual-unicast.

In yet another embodiment, a session (e.g., RTP or RTCP) packet for transmission on a network, comprises:

- a source network address of a first participant to a session;
- a destination network address associated with a session monitor;
- a network address of a second participant to the session; and
- session information associated with the session.

In RTCP applications, the session packet can further include a first session id associated with the first participant and a second session id associated with the second participant. The network address of the second participant is typically located in the application or APP field of the RTCP packet.

In other embodiments, the systems and hardware are provided to implement the above embodiments.

The various embodiments of the present invention can have numerous advantages. For example, the window of opportunity for confusing concurrent sessions and attributing data in RTCP packets to the wrong session can be much smaller than with current architectures. The window of opportunity for possible confusion using the above algorithm(s) exists only when two different endpoints join different sessions at the same time and with the

same SSRC ids. This window of opportunity or startup interval closes once either of the endpoints (or their peers) has sent an RTCP packet with a reception block corresponding to either endpoint. Once the reception block is exchanged, the SSRC ids of both parties to the session are known to the monitor. Before such an exchange, the monitor typically has only the SSRC id and network address of one party to the session. The SSRC id and network address of the other party is unknown. The startup interval is typically fairly short, *e.g.*, typically on the order of 5 seconds or less. Even this potential period of confusion can be eliminated by choosing an SSRC id for each endpoint that is globally unique. The use of the active session table and network address to define the session (rather than only pairings of SSRC ids) can, after the startup interval, at least substantially eliminate misinterpretation of RTCP packets and incorrect analysis of performance data. The accuracy of the algorithm(s) in matching RTCP packets with the corresponding session results in more accurate statistical analysis of the communication link in the network.

The above-described embodiments and configurations are neither complete nor exhaustive. As will be appreciated, other embodiments of the invention are possible utilizing, alone or in combination, one or more of the features set forth above or described in detail below.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a session;

Fig. 2 is a flowchart depicting an algorithm according to an embodiment of the present invention;

Fig. 3 is a block diagram of a monitor according to an embodiment of the present invention; and

Fig. 4 is a flowchart depicting an algorithm according to an embodiment of the present invention.

5

DETAILED DESCRIPTION

The Monitor

10 The monitor of the present invention, in one embodiment, includes data structures and applications to track multiple concurrent sessions. Referring to Fig. 3, the monitor 300 includes in memory 302 an orphan table 304 containing, *inter alia*, a listing of unmatched session endpoints or orphans, an active session table 308 containing, *inter alia*, a listing of matched session endpoints, a parser 304 to parse RTCP packets and locate selected fields, and a matcher 308 to compare selected fields in the packet with selected fields in the orphan and active session tables 304, 308.

15

The orphan table 304 typically includes, for each known endpoint, transport address of a first endpoint in a selected session, SSRC of the first endpoint, optionally SSRC of a second endpoint in the selected session, and associated data structures in the RTCP packet, such as jitter, packet loss, and round-trip time, related to the selected session. The network address of the second endpoint is unknown. The SSRC of the second (unknown) endpoint is typically contained in a reception report of the RTCP packet. The SSRC of the second endpoint is generally unavailable in the first RTCP packet received by the monitor from an endpoint in a session. The SSRC of the second endpoint is commonly generated only after

20

one session participant receives an initial packet from the other session packet (or after packets have been exchanged by the session participants).

The active session table 308 typically includes, for each known or matched session, transport address of a first endpoint in a selected session, optionally SSRC of the first endpoint, network address of a second endpoint in the selected session, optionally SSRC of the second endpoint, and associated data structures in the RTCP packet, such as jitter, packet loss, and round-trip time, associated with the selected session. As will be appreciated, the entry in the orphan table in which the network address of the first endpoint is known is moved to the active session table when the monitor is able to identify the network address of the second (unknown) endpoint. As discussed below, this is typically determined by matching information in a newly received RTCP packet with one or more entries in the orphan table.

Operations of the Matching Algorithms

Sub A3> The operations of the matching algorithm(s) in the monitor 300 will now be discussed with reference to Fig. 2. Referring to Fig. 2, a packet is received by the monitor in step 200. Parser 312 parses the packet to locate selected fields, which typically are the source transport address, source SSRC id ("the endpoint SSRC id"), if present the destination transport address of the other session participant (which is possibly in the application APP field), and, if present, the destination SSRC id of the other session participant in the receiver report blocks (the SSRC's in the receiver report blocks are hereinafter referred to as the "reception report SSRC id"). As will be appreciated, the reception report is typically a report regarding

Sub A³

the characteristics of the communication link, such as the condition of the voice stream experienced since the last reception report.

In step 204, the matcher 316 first determines whether the APP field includes the destination transport address. If the APP field contains the destination transport address, the
5 matcher 316 in step 208 updates the orphan and active tables. The orphan table can contain an entry corresponding to a first session participant where the first participant is not configured to place the transport address of the other (second) party in the APP field. In other configurations, however, the first packet received is from a party that is configured to include the destination transport address of the other session participant in the APP field and
10 no entry will be made in the orphan table as the packet itself contains the transport addresses of both session participants. In any event, the matcher 316 removes the entry, if any, from the orphan table and, if necessary, creates a new entry in the active session table for the session. If an entry is already in the active session table, the entry is updated in step 208 by updating the associated data fields with the performance information included in the packet.
15 The monitor 300 then returns to step 300 to await the next RTCP packet.

If the APP field does not include a transport address, the monitor 300 next proceeds to step 212 where the matcher 316 determines whether the source identified in the packet corresponds to an endpoint in the active session table. Matcher 316 makes this determination by matching on transport address or UDP and endpoint SSRC pair.

20 If the matcher 316 receives a hit (or match), the matcher 316, in step 216, updates the fields in the active session table and returns to step 200. As noted, each entry in the active

session table includes at least endpoint pairs (identified by transport address, UDP and/or endpoint SSRC's).

If the matcher 316 receives a no hit (or no match), the matcher 316, in step 220, determines if the source or endpoint in the packet has a matching entry in the orphan table.

5 This is typically determined by matching on transport address or UDP and endpoint SSRC pair.

Sub A⁴ If the matcher 316 receives a hit, the monitor in step 224 updates the entry for the orphan session. This is typically done by updating the other party's SSRC id (if available) and updating the associated data in the packet. As noted, each entry in the orphan session table
10 includes at least UDP or transport address of an endpoint, an endpoint SSRC id, and optionally reception report SSRC id.

If the matcher 316 receives a no hit, the matcher 316 in step 228 creates an entry (if necessary) in the orphan session table 304.

15 After both steps 224 and 228, the monitor 300 next determines in step 232 whether a pair of entries corresponding to a pair of orphans have matching endpoint SSRC and reception report SSRC's. If no match is found, the monitor proceeds to step 200 to await the next packet. If a match is found, the monitor in step removes the matched end-point entries from the orphan table and creates one entry in the active session table with the data in the two entries in corresponding fields. The monitor then returns to step 200 to await the next packet.

20 Fig. 4 depicts the algorithm for a computational component in the first endpoint that is configured to input another (second) endpoint's SSRC id into the APP field. In step 400, the first endpoint receives an RTCP packet from the second endpoint participating in a session

-11-

algorithm(s) can be used with a unicast architecture where a packet redirector or copier such as a sniffer, probe, or router is positioned between the two parties to the session to intercept RTCP packets and either send a copy of the packet or the packet itself to the monitor. In yet another embodiment, the orphan and active session tables can be combined into one table, with a flag or other indicator being used to indicate when an entry has been completed (*i.e.*, both endpoints to the session have been identified by transport address).

The present invention, in various embodiments, includes components, methods, processes, systems and/or apparatus substantially as depicted and described herein, including various embodiments, subcombinations, and subsets thereof. Those of skill in the art will understand how to make and use the present invention after understanding the present disclosure. The present invention, in various embodiments, includes providing devices and processes in the absence of items not depicted and/or described herein or in various embodiments hereof, including in the absence of such items as may have been used in previous devices or processes, e.g. for improving performance, achieving ease and/or reducing cost of implementation.

The foregoing discussion of the invention has been presented for purposes of illustration and description. The foregoing is not intended to limit the invention to the form or forms disclosed herein. Although the description of the invention has included description of one or more embodiments and certain variations and modifications, other variations and modifications are within the scope of the invention, e.g. as may be within the skill and knowledge of those in the art, after understanding the present disclosure. It is intended to obtain rights which include alternative embodiments to the extent permitted, including

alternate, interchangeable and/or equivalent structures, functions, ranges or steps to those claimed, whether or not such alternate, interchangeable and/or equivalent structures, functions, ranges or steps are disclosed herein, and without intending to publicly dedicate any patentable subject matter.